Throughput-Optimal Load Balancing for Intra Datacenter Networks

Sucha Supittayapornpong, Michael J. Neely
University of Southern California
Email: supittay@usc.edu, mjneely@usc.edu

arXiv:1612.01684v1 [cs.NI] 6 Dec 2016

***Abstract*—Traffic load-balancing in datacenters alleviates hot spots and improves network utilization. In this paper, a stable in-network load-balancing algorithm is developed in the setting of software-defined networking. A control plane configures a data plane over successive intervals of time. While the MaxWeight algorithm can be applied in this setting and offers certain throughput optimality properties, its *bang-bang* control structure rewards single flows on each interval and prohibits link-capacity sharing. This paper develops a new algorithm that is throughput-optimal and allows link-capacity sharing, leading to low queue occupancy. The algorithm deliberately imitates weighted fair queueing, which provides fairness and graceful interaction with TCP traffic. Inspired by insights from the analysis, a heuristic improvement is also developed to operate with practical switches and TCP flows. Simulations from a network simulator shows that the algorithm outperforms the widely-used equal-cost multipath (ECMP) technique.**

I. Introduction

Datacenter networks serve as infrastructure for search engines, social networks, cloud computing, etc. Due to potentially high traffic loads, load-balancing becomes an important solution to improve network utilization and alleviate hot spots [1]–[4]. A widely-used technique is equal-cost multipath (ECMP), where traffic flows are split equally according to the number of available equal-cost next-hops. However, ECMP does not take into account actual traffic and is susceptible to asymmetric topology [5]–[7]. Further, the deployment of ECMP is limited due to its equal-cost constraint [8].

Traffic load-balancing can be implemented using software-defined networking (SDN). An SDN switch, a network device supporting layer-2 and layer-3 operations in OSI architecture, consists of a data plane and control plane [9].[1] The data plane forwards packets according to given rules and operates at a fast timescale, e.g., 1ns. The control plane sets those rules and operates at a much slower timescale, e.g., 1ms. Several traffic load-balancing algorithms can be implemented through the control plane reconfiguration.

Existing traffic load-balancing methods for datacenter networks distribute traffic according to network capacity and measured traffic. Weighted-cost multipath [7] distributes traffic according to path capacity. Centralized algorithms, such as Hedera [10] and Niagara [5], take advantage of global traffic information to split traffic at a coarse timescale. Load-balancing with a finer timescale has been implemented in DeTail [11] and CONGA [6] using the *in-network* technique, where decisions are made at switches inside a network without any central controller. Conceptually, these approaches attempt to distribute traffic over available network resources. However, the path-based approach in CONGA limits scalability, and the packet-by-packet dispersion in DeTail needs TCP with out-of-order resilience. Further, these algorithms do not come with analytical optimality proofs, and it is not clear if they are *throughput optimal*.

[1]An SDN switch should not be confused with a crossbar switch, which may reside in the data plane of an SDN switch.

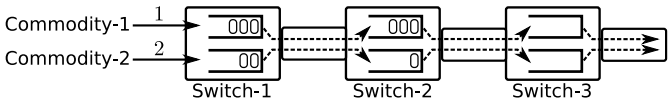

Fig. 1. MaxWeight example: Let w_{ij}^d be the weight of commodity d over the link from switch i to switch j. All weights in this figure are $\left(w_{12}^1, w_{12}^2, w_{23}^1, w_{23}^2\right) = (0, 1, 3, 1)$.

An algorithm is throughput optimal if it stably supports any feasible traffic load, so that average backlog is bounded [12]. Specifically, a throughput-optimal algorithm utilizes the entire network capacity and can distribute traffic to any portion of the network to maintain network stability. MaxWeight [12] is a well-known throughput-optimal algorithm and has been studied for packet radio [12], switching [13], and inter-datacenter networking [14]. It has been generalized to optimize power allocation [15], throughput [16], etc. Practical aspects of MaxWeight such as finite buffer capacity and fairness with TCP connections have been studied in [17]–[19]. However, MaxWeight is not suitable for in-network load-balancing because it prohibits the sharing of link capacity at the data plane's timescale and thus causes high queue occupancy. This queue size has a finite average, but the size of the longer timescale makes that average unacceptably large.

The MaxWeight algorithm is illustrated by the example in Fig. 1. Two traffic commodities share three links passing through switches $1, 2$ and 3. Time is slotted. The slot size equals the length of the decision update interval (the control plane's timescale). The capacity of each link is 3 packets per slot. Every switch has a dedicated queue for each commodity. In every time slot, MaxWeight calculates, for each link and commodity, a weight equal to the differential backlog between a queue and its next-hop queue. For that slot, the entire capacity of the link is allocated to the commodity with the maximum non-negative weight, while a commodity with negative weight is ignored. For example, commodity 2 is served on the link between switches 1 and 2 in Fig. 1, and commodity 1 is served

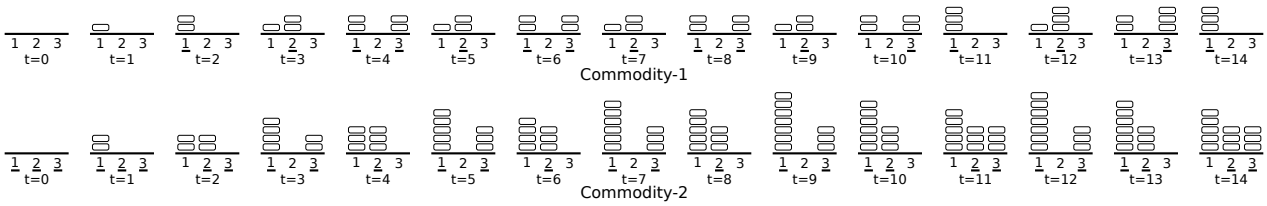

Fig. 2. Timeline of queue occupancy under MaxWeight: a small box represents a packet in a queue. Switch i is represented by the number $i \in \{1, 2, 3\}$ under the long line. The short line under the number indicates that the commodity at the numbered switch is served in that particular time slot. The occupancy pattern repeats after $t = 14$, which is similar to the pattern at $t = 11$.

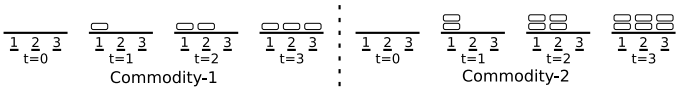

Fig. 3. Timeline of queue occupancy under the ideal algorithm

on the link between switches 2 and 3.

Let the arrival rates to switch 1 of commodities 1 and 2 be respectively 1 and 2 packets per slot. The timeline of queue evolution is shown in Fig. 2. MaxWeight is effective and always transmits three packets per slot after $t = 11$. This effectiveness requires a sufficient amount of queue backlog. For example, commodity 2 backlog at switch 1 is always at least 5 for times $t \geq 11$. This occupancy might be acceptable. However, if the control plane is reconfigured at a slow timescale relative to the link capacity, the queue occupancy can be very high. For example, with a 1ms update interval, 10Gbps link speed, and 1kB packet size, each link can serve 1250 packets per slot (rather than just 3). This multiplies queue backlog in the timeline of Fig. 2 by a factor $1250/3$, so the minimum queue backlog of commodity 2 at switch 1 is $5 \times (1250/3) \approx 2083$ for $t \geq 11$. Another undesirable property of MaxWeight is that queue occupancy scales linearly with the number hops, as shown in [20], [21]. In fact, Fig. 2 is inspired by an example in [20].

In practice, the MaxWeight mechanism with a long update interval leads to i) large buffer memory, ii) packet drops, iii) high latency, and iv) burstiness (no capacity sharing during an update interval). Issues (i)–(iii) can be alleviated partially by the techniques in [21]–[24]. However, issue (iv) resides in the decision making mechanism of MaxWeight and still persists under those techniques. The situation is worse when issues (ii) and (iv) interact with TCP congestion control, causing slow flow rate and under utilization. To put it into theoretical perspective, even though MaxWeight solves a network stability problem with $O(1)$ average queue size, the constant factor is too large for a practical system with a long update interval. Note that an ideal algorithm for the example in Fig. 1 always serves 1 and 2 packets of commodities 1 and 2 by sharing the link capacity as shown in Fig. 3.

In this paper, a new throughput-optimal algorithm is developed. The algorithm shares link capacity among commodities during an update interval, resulting in low queue occupancy and low latency. The key challenge is to design a model and an algorithm that are analyzable, provably optimal, and practically implementable at the same time. The algorithm imitates the weighted fair queueing (WFQ) [25], [26] available in practical switches to provide fairness and low latency among TCP flows in practice. A general intra data center network may have an exponential number of paths, and our algorithm comes with an optimality proof considering all possible paths using per-commodity queueing which grows linearly with the number of commodities. This is also a key distinct aspect from the path-based algorithm in [27].

Section III develops the throughput-optimal algorithm. Inspired by this algorithm, Section IV presents an enhanced algorithm that includes heuristics to cope with practical aspects, including queue information dissemination, queue approximation, and packet reordering issues in TCP. This heuristic algorithm uses local queue information and local measured traffic to hash TCP flows to next-hop switches and set weights of the weighted fair queueing. The hash-based mechanism is chosen to reduce packet reordering, which is not possible for DeTail [11]. Simulation results in Section V show that both proposed algorithms outperform MaxWeight and ECMP algorithms in ideal simulation and in more realistic simulation with OMNeT++ [28].

II. System Model and Design

The fast timescale of the data plane operates over slotted time $t \in \mathbb{Z}_+$, where $\mathbb{Z}_+ = \{0, 1, 2, \ldots\}$. The control plane configures the data plane every T slots, where T is a positive integer. Thus, reconfigurations occur at times in the set $\mathcal{T} = \{0, T, 2T, \ldots\}$.

A. Topology and Routing

An intra datacenter network is the interconnection of switches and destinations (such as servers), as shown in Fig. 4. Traffic designated for a particular destination d is called *commodity* d traffic. Let $\mathcal{S}$ be the set of all switches and $\mathcal{D}$ be the set of all destinations (commodities). A link between switches i and j is bi-directional with capacity c_{ij} from i to j, and capacity c_{ji} in the reverse direction. Define $c_{ij} = 0$ if link (i, j) does not exist or if $i = j$.

Each switch must decide where to send its packets next. Define $\mathcal{H}_i^d \subseteq \mathcal{S}$ as the set of next-hop switches available to commodity d packets in switch i, for all $i \in \mathcal{S}$ and $d \in \mathcal{D}$. In practice, these sets can be obtained from manual configuration or from other routing mechanisms. Define the set of all next-hop switches from switch i as $\mathcal{H}_i = \cup_{d \in \mathcal{D}} \mathcal{H}_i^d$, and the set of previous-hop switches as $\mathcal{P}_i^d = \left\{ j \in \mathcal{S} : i \in \mathcal{H}_j^d \right\}$ for $i \in \mathcal{S}, d \in \mathcal{D}$. These sets are illustrated in Fig. 5. Define $\mathcal{D}_{ij} = \left\{ d \in \mathcal{D} : j \in \mathcal{H}_i^d \right\}$ as the set of all commodities utilizing a link

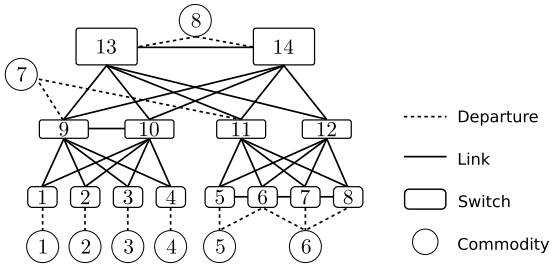

Fig. 4. An example network with $\mathcal{S} = \{1, 2, \ldots, 14\}$ and $\mathcal{D} = \{1, 2, \ldots, 8\}$

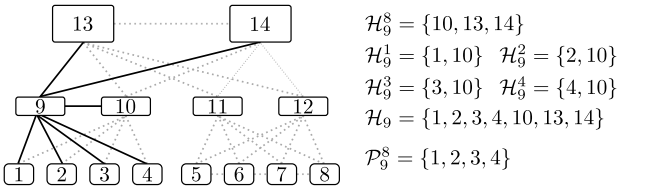

Fig. 5. Example of sets of switches at switch 9. Note that $\mathcal{P}_9^8$ must not contain 10 to avoid loops and imposes that $9 \notin \mathcal{H}_{10}^8$.

from switch i to switch j for $i, j \in \mathcal{S}$. Note that this model allows arbitrary path lengths, and can be applied to existing topologies in [1]–[4], [8].

B. Traffic

Switch i receives $a_i^d(t)$ commodity-d packets from external sources at time t (for $i \in \mathcal{S}, d \in \mathcal{D}$). The external source represents a group of servers or a link connecting to the outside of the network. For each $i \in \mathcal{S}$, $d \in \mathcal{D}$, the arrival process $\{a_i^d(t)\}_{t=0}^{\infty}$ is independent and identically distributed (i.i.d.) across time slots. The i.i.d. assumption is useful for a simple and insightful analysis. The resulting algorithm developed under this assumption inspires a heuristic algorithm in Section IV that does not require i.i.d. arrivals and works gracefully with TCP traffic.

Recall that c_{ij} is the capacity of the link between switch i and switch j, for $i, j \in \mathcal{S}$. Let b_i^d be the capacity of the link between switch $i \in \mathcal{S}$ and destination $d \in \mathcal{D}$. Set $b_i^d = 0$ if switch i does not have a direct link to destination d. Assume arrivals and link capacities are always bounded by a constant $\delta > 0$, so $0 \le c_{ij} \le \delta$, $0 \le a_i^d(t) \le \delta$, $0 \le b_i^d \le \delta$ for all $i, j \in \mathcal{S}, d \in \mathcal{D}, t \in \mathbb{Z}_+$.

C. Decision Variables

Decision variables are defined for every link connecting switch i to its next-hop switch j, for $i \in \mathcal{S}, j \in \mathcal{H}_i$. Recall that $\mathcal{D}_{ij}$ is a set of commodities using the link. At configuration time $t \in \mathcal{T}$, the control plane in switch i chooses a *control plane decision variable* $x_{ij}^d(t, T)$ for $d \in \mathcal{D}_{ij}$, which represents a constant transmission rate allocated to commodity d (in units of packets) for the entire T-slot interval. Define $x_{ij}^d(t, T) = 0$ for $d \in \mathcal{D} \backslash \mathcal{D}_{ij}$. The control plane decisions for link (i, j) are chosen to satisfy the link capacity constraint:

$$\sum_{d \in \mathcal{D}_{ij}} x_{ij}^d(t, T) \le T c_{ij}.$$

Once $x_{ij}^d(t, T)$ is determined, no more than $x_{ij}^d(t, T)$ commodity-d packets can be transmitted by the data plane during an interval $\{t, \ldots, t + T - 1\}$. For example, the data plane can impose a token bucket mechanism. Let $x_{ij}^d(t)$ be *data plane decision variable* that represents the transmission rate assigned by the data plane to commodity d on link (i, j) for slot t. These are chosen to satisfy

$$x_{ij}^d(t, T) = \sum_{\tau=t}^{t+T-1} x_{ij}^d(\tau) \quad \text{for } i \in \mathcal{S}, j \in \mathcal{H}_i, t \in \mathcal{T} \quad (1)$$

$$\sum_{d \in \mathcal{D}_{ij}} x_{ij}^d(t) \le c_{ij} \quad \text{for } i \in \mathcal{S}, j \in \mathcal{H}_i, t \in \mathbb{Z}_+. \quad (2)$$

D. Queues

Packets are queued at each switch according to their commodity. Let $Q_i^d(t)$ be the number of commodity-d packets queued at switch i on slot t. The value $Q_i^d(t)$ is also called the *commodity-d backlog* and satisfies:

$$Q_i^d(t+1) \le \left[Q_i^d(t) - \sum_{j \in \mathcal{H}_i^d} x_{ij}^d(t) - b_i^d \right]_+ + \sum_{j \in \mathcal{P}_i^d} x_{ji}^d(t) + a_i^d(t) \quad \text{for } i \in \mathcal{S}, d \in \mathcal{D}, \quad (3)$$

where $[z]_+ = \max[0, z]$. Note that $\sum_{j \in \mathcal{H}_i^d} x_{ij}^d(t)$ denotes output transmission rate to next-hop switches, and $\sum_{j \in \mathcal{P}_i^d} x_{ji}^d(t)$ denotes receiving transmission rate from previous-hop switches. Inequality (3) is an inequality rather than an equality because the actual amount of new endogenous arrivals on slot t may be less than $\sum_{j \in \mathcal{P}_i^d} x_{ji}^d(t)$ if previous switches j do not have enough commodity d backlog to fill the assigned transmission rate $x_{ji}^d(t)$. It can be shown that the backlog at time $t \in \mathcal{T}$ satisfies for every $i \in \mathcal{S}, d \in \mathcal{D}$:

$$Q_i^d(t+T) \le \left[Q_i^d(t) - \sum_{j \in \mathcal{H}_i^d} x_{ij}^d(t, T) - T b_i^d \right]_+ + \sum_{j \in \mathcal{P}_i^d} x_{ji}^d(t, T) + \sum_{\tau=t}^{t+T-1} a_i^d(\tau). \quad (4)$$

Note that, while a common queue for each commodity is not available in practical switches, it can be heuristically implemented by queues in an SDN switch in Section IV.

E. Stability and Assumption

Definition 1 (Queue Stability [22])**.** *A queue with backlog $\{Z(t) \ge 0 : t \in \mathbb{Z}_+\}$ is strongly stable if*

$$\limsup_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\left[Z(\tau)\right] < \infty.$$

Definition 2 (Network Stability [22])**.** *A network is strongly stable when every queue in the network is strongly stable.*

The arrival and departure rates are assumed to satisfy a standard Slater condition:

Assumption 1 (Slater Condition)**.** *There exists an* $\epsilon > 0$ *and a randomized policy* $\left\{x_{ij}^{d*}(t)\right\}_{i\in\mathcal{S},d\in\mathcal{D},j\in\mathcal{H}_i^d}$ *with*

$$\mathbb{E}\left[\sum_{j\in\mathcal{P}_i^d} x_{ji}^{d*}(t) + a_i^d(t) - \sum_{j\in\mathcal{H}_i^d} x_{ij}^{d*}(t) - b_i^d\right] < -\epsilon \quad \text{for all } i\in\mathcal{S}, d\in\mathcal{D}, t\in\mathbb{Z}_+,$$

and the randomized policy satisfies the constraint (2)*.*

Note that Assumption 1 is stated in terms of only the data plane decision variables $x_{ij}^{d*}(t)$ and the constraint (2). While the control plane decisions are used in the algorithm and the additional constraint (1) is satisfied by the algorithm, those are not used in the Slater condition.

III. Throughput-Optimal Algorithm

A. The Algorithm

The novel throughput-optimal algorithm runs distributively at every switch. Let K be a positive real number, and $x_{ij}^d(-T,T) = 0$ for all $i,j\in\mathcal{S}, d\in\mathcal{D}$. At every reconfiguration time $t\in\mathcal{T}$, switch i executes Algorithm 1 for each link connecting to a next-hop switch j for $i\in\mathcal{S}, j\in\mathcal{H}_i$. Recall that $x_{ij}^d(t,T) = 0$ for $d\in\mathcal{D}\backslash\mathcal{D}_{ij}$. The function $\mathrm{round}\,[z]$ rounds the real-valued z to its closest integer.

```
// switch i, link to switch j, time t //
y_ij^d(t) ← Q_i^d(t) − Q_j^d(t) + x_ij^d(t−T,T) for d ∈ D_ij
k_ij(t) ← max[1, min[K, (1/(Tc_ij)) Σ_{d∈D_ij} [y_ij^d(t)]_+]]
if Σ_{d∈D_ij} round[[y_ij^d(t)]_+ / k_ij(t)] ≤ Tc_ij then
    x_ij^d(t,T) ← round[[y_ij^d(t)]_+ / k_ij(t)] for d ∈ D_ij
    return {x_ij^d(t,T)}_{d∈D_ij}
else
    return result of Algorithm 2
end if
```

Algorithm 1: Throughput-optimal rate allocation

```
// switch i, link to switch j, time t //
y_ij^d(t) ← Q_i^d(t) − Q_j^d(t) + x_ij^d(t−T,T) for d ∈ D_ij
v_ij^d ← 0 for d ∈ D_ij
for n = 1 to Tc_ij do
    d_n ← argmax_{e∈D_ij} [y_ij^e(t) − v_ij^e k_ij(t)]
    v_ij^{d_n} ← v_ij^{d_n} + 1
end for
x_ij^d(t,T) ← v_ij^d for d ∈ D_ij
return {x_ij^d(t,T)}_{d∈D_ij}
```

Algorithm 2: Packet-filling algorithm (unaccelerated version)

B. Intuitions

Algorithm 1 solves problem (5) with the value of $k_{ij}(t)$ that depends on local queue information and previous decisions. This $k_{ij}(t)$ is deliberately introduced, just as a solution of problem (5) imitates weighted fair queueing, which provides fairness and low latency in practice [25], [26].

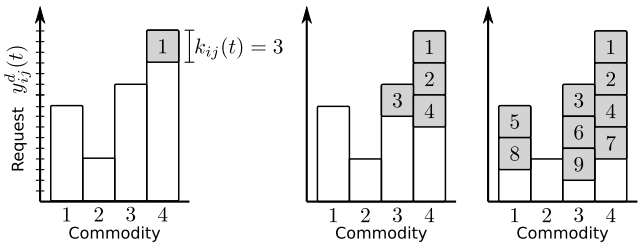

Fig. 6. Packet-filling Algorithm 2 iteratively fulfills the requests. An iteration number is indicated in a gray box. In this example, $Tc_{ij} = 9$ and the first iteration (the plot on the left) allocates rate to commodity 4. The algorithm allocates $2, 0, 3, 4$ rates to commodities $1, 2, 3, 4$.

$$\begin{aligned}
\text{Minimize} \quad & \sum_{d\in\mathcal{D}_{ij}} \left\{ x_{ij}^d(t,T)\left[Q_j^d(t) - Q_i^d(t)\right] + \right. \\
& \left. \frac{k_{ij}(t)}{2}\left[x_{ij}^d(t,T) - \frac{x_{ij}^d(t-T,T)}{k_{ij}(t)}\right]^2 \right\} \qquad (5) \\
\text{Subject to} \quad & \sum_{d\in\mathcal{D}_{ij}} x_{ij}^d(t,T) \le Tc_{ij} \\
& x_{ij}^d(t,T) \in \mathbb{Z}_+ \qquad \text{for all} \quad d\in\mathcal{D}_{ij}
\end{aligned}$$

In Algorithm 1, $y_{ij}^d(t)$ represents a request for transmission rate of commodity d. It also indicates how well the previous rates were allocated. Under allocation of $x_{ij}^d(t-T,T)$ increases queue backlog $Q_i^d(t)$, which tends to increase the request $y_{ij}^d(t)$. It is easy to see that, if the queue backlogs $Q_i^d(t)$ and $Q_j^d(t)$ are about the same, then the request is about the same as its previous value. This behavior is smoother than MaxWeight.

The requests are fulfilled in two situations. i) When the total requests are roughly within KTc_{ij}, i.e., $\sum_{d\in\mathcal{D}_{ij}} \mathrm{round}\left[\left[y_{ij}^d(t)\right]_+ / k_{ij}(t)\right] \le Tc_{ij}$, the requests are fulfilled in WFQ fashion, which can be seen by considering $k_{ij}(t) = \frac{1}{Tc_{ij}}\sum_{d\in\mathcal{D}_{ij}}\left[y_{ij}^d(t)\right]_+$ and

$$\begin{aligned}
x_{ij}^d(t,T) &= \mathrm{round}\left[\left[y_{ij}^d(t)\right]_+ / k_{ij}(t)\right] \\
&= \mathrm{round}\left[\frac{\left[y_{ij}^d(t)\right]_+}{\sum_{e\in\mathcal{D}_{ij}}\left[y_{ij}^e(t)\right]_+} \times Tc_{ij}\right]. \qquad (6)
\end{aligned}$$

ii) The other is an extreme situation for stability analysis, which occurs when a network operates near its capacity, which may not be the case in practice due to TCP congestion control. This case is solved by Algorithm 2[2], as illustrated by Fig. 6. It is easy to see the fairness introduced by $k_{ij}(t)$. Without $k_{ij}(t)$, i.e., $k_{ij}(t)$ is always 1, the allocation in Fig. 6 will be $0, 0, 2, 7$ for commodities 1 to 4, which may cause a fairness issue with TCP flows [19].

[2]Algorithm 2 can be accelerated by fulfilling multiple requests per iteration. For example, the requests in Fig. 6 can be fulfilled in three iterations.

C. Correctness of Algorithm 1

This subsection shows that Algorithm 1 returns an optimal solution of problem (5). To simplify notation in this section, the time index of variables and constants in problem (5) are omitted. Let z_{ij}^d denote $x_{ij}^d(t-T,T)$.

When commodity d is allocated an integer service value v_{ij}^d, define its contribution to the cost function of problem (5) as

$$g_{ij}^d(v_{ij}^d) = v_{ij}^d\big[Q_j^d - Q_i^d\big] + \frac{k_{ij}}{2}\left[v_{ij}^d - \frac{z_{ij}^d}{k_{ij}}\right]^2.$$

The cost difference of getting another service allocation is

$$g_{ij}^d(v_{ij}^d+1) - g_{ij}^d(v_{ij}^d) = -\left[Q_i^d - Q_j^d + z_{ij}^d - \frac{k_{ij}}{2} - k_{ij}v_{ij}^d\right]. \tag{7}$$

Since the cost function in problem (5) is minimized, commodity d only accepts a transmission allocation if $g_{ij}^d(v_{ij}^d+1) \le g_{ij}^d(v_{ij}^d)$. The cost difference in (7) is monototically increasing in v_{ij}^d. Therefore, commodity d receives transmission allocation at most:

$$\begin{aligned} x_{ij}^{d(\max)} &= \min\big\{v \in \mathbb{Z}_+ : g_{ij}^d(v+1) - g_{ij}^d(v) > 0\big\} &(8)\\ &= \min\left\{v \in \mathbb{Z}_+ : \frac{Q_i^d - Q_j^d + z_{ij}^d}{k_{ij}} - 0.5 < v\right\}\\ &= \text{round}\left[\big[Q_i^d - Q_j^d + z_{ij}^d\big]_+/k_{ij}\right]. &(9)\end{aligned}$$

Lemma 1. *When $k_{ij} > 0$ and $\sum_{d\in\mathcal{D}_{ij}} x_{ij}^{d(max)} \le Tc_{ij}$, the optimal solution of problem* (5) *is $x_{ij}^d = x_{ij}^{d(max)}$ for $d \in \mathcal{D}_{ij}$.*

Proof. For any $v_{ij}^d \in \left\{0,1,\dots,x_{ij}^{d(\max)}\right\}, d \in \mathcal{D}_{ij}$, the given implies that $\sum_{d\in\mathcal{D}_{ij}} v_{ij}^d \le \sum_{d\in\mathcal{D}_{ij}} x_{ij}^{d(\max)} \le Tc_{ij}$. So, any chosen v_{ij}^d in $\left\{0,1,\dots,x_{ij}^{d(\max)}\right\}$ leads to a feasible solution of problem (5). Note that the objective function of problem (5) is separable, $\sum_{d\in\mathcal{D}_{ij}} g_{ij}^d(v_{ij}^d)$, and is minimized. The definition of $x_{ij}^{d(\max)}$ in (8) implies $g_{ij}^d(v) > g_{ij}^d(x_{ij}^{d(\max)})$ for any $v > x_{ij}^{d(\max)}$, so any $v > x_{ij}^{d(\max)}$ is not optimal. Thus, i) if $x_{ij}^{d(\max)} = 0$, $x_{ij}^d = 0 = x_{ij}^{d(\max)}$ minimizes problem (5) with respect to commodity d. ii) If $x_{ij}^{d(\max)} > 0$, the definition of $x_{ij}^{d(\max)}$ in (8) implies $g_{ij}^d(x_{ij}^{d(\max)}) \le g_{ij}^d(v)$ for any $v \in \left\{0,1,\dots,x_{ij}^{d(\max)}-1\right\}$, so $x_{ij}^d = x_{ij}^{d(\max)}$ minimizes the problem with respect to commodity d. □

Lemma 1 implies that Algorithm 1 solves problem (5) when the first (if-)condition is met. The other case of Algorithm 1 can be shown by the following property. When commodity d is allocated v_{ij}^d transmission rate, define the unfulfilled level of commodity-d's request as

$$l_{ij}^d(v_{ij}^d) = Q_i^d - Q_j^d + z_{ij}^d - k_{ij}v_{ij}^d. \tag{10}$$

Lemma 2. *When $k_{ij} > 0$, for any commodities $d, e \in \mathcal{D}_{ij}$ whose allocated rates are respectively v_{ij}^d and v_{ij}^e, the following holds:*

i) If $l_{ij}^d(v_{ij}^d) = l_{ij}^e(v_{ij}^e)$, then

$$g_{ij}^d(v_{ij}^d+1) + g_{ij}^e(v_{ij}^e) = g_{ij}^d(v_{ij}^d) + g_{ij}^e(v_{ij}^e+1).$$

ii) If $l_{ij}^d(v_{ij}^d) > l_{ij}^e(v_{ij}^e)$, then

$$g_{ij}^d(v_{ij}^d+1) + g_{ij}^e(v_{ij}^e) < g_{ij}^d(v_{ij}^d) + g_{ij}^e(v_{ij}^e+1).$$

Proof. It holds from equations (7) and (10) that

$$\begin{aligned} &g_{ij}^d(v_{ij}^d+1) - g_{ij}^d(v_{ij}^d) - g_{ij}^e(v_{ij}^e+1) + g_{ij}^e(v_{ij}^e)\\ &\quad= -\left[l_{ij}^d(v_{ij}^d) - \frac{k_{ij}}{2}\right] + \left[l_{ij}^e(v_{ij}^e) - \frac{k_{ij}}{2}\right]\\ &\quad= -l_{ij}^d(v_{ij}^d) + l_{ij}^e(v_{ij}^e). &(11)\end{aligned}$$

In case (i), $l_{ij}^d(v_{ij}^d) = l_{ij}^e(v_{ij}^e)$ implies that $g_{ij}^d(v_{ij}^d+1) - g_{ij}^d(v_{ij}^d) - g_{ij}^e(v_{ij}^e+1) + g_{ij}^e(v_{ij}^e) = 0$, which proves (i). Case (ii) can be proven similarly by substituting $-l_{ij}^d(v_{ij}^d) + l_{ij}^e(v_{ij}^e) < 0$ into equation (11) and rearranging terms. □

Lemma 2 implies that allocating rate to the commodity with the highest unfulfilled level reduces the total objective the most. Specifically, let v_{ij}^d be the current rate allocation of commodity d. For $d^* = \text{argmax}_{d\in\mathcal{D}_{ij}} l_{ij}^d(v_{ij}^d)$, Lemma 2 implies that

$$\begin{aligned} &g_{ij}^{d^*}(v_{ij}^{d^*}+1) + \sum_{d\in\mathcal{D}_{ij}\setminus\{d^*\}} g_{ij}^d(v_{ij}^d) \le\\ &\qquad g_{ij}^e(v_{ij}^e+1) + \sum_{d\in\mathcal{D}_{ij}\setminus\{e\}} g_{ij}^d(v_{ij}^d) \quad \text{for all } e \in \mathcal{D}_{ij}.\end{aligned}$$

The above property ensures that the iterative allocation in Algorithm 2 greedily optimizes problem (5) when event $\sum_{d\in\mathcal{D}_{ij}} \text{round}\left[\big[y_{ij}^d(t)\big]_+/k_{ij}(t)\right] > Tc_{ij}$ occurs. It can be proven by contradiction that commodity d gets at most $x_{ij}^{d(\max)}$ rate for all $d \in \mathcal{D}_{ij}$. Then, the algorithm always allocates the entire transmission rate, as an implication of the event. Therefore, Algorithm 2 returns an optimal solution of problem (5).

Theorem 1. *Given $K > 0$, Algorithm 1 solves problem 5 with $k_{ij}(t) \in [1,K]$, where $k_{ij}(t)$ is defined in the algorithm.*

Proof. The theorem is the consequence of Lemmas 1 and 2 and the fact that $k_{ij}(t) \in [1,K]$. □

D. Stability Analysis

Problem 5 with $k_{ij}(t) \in [1,K]$ is shown to be a class of throughput-optimal policies. Let $Q(t) = \big(Q_i^d(t)\big)_{i\in\mathcal{S},d\in\mathcal{D}}$ be a vector of all backlogs at time t. Define $\|z\|_1$ as the l_1-norm, e.g., $\|Q(t)\|_1 = \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}} Q_i^d(t)$.

Theorem 2. *When Assumption 1 holds, the network is strongly stable:*

$$\limsup_{N\to\infty} \frac{1}{NT}\sum_{n=0}^{N-1}\sum_{\tau=0}^{T-1} \mathbb{E}\left[\|Q(nT+\tau)\|_1\right] \le \frac{G_1}{\epsilon} + G_2,$$

where $G_1 = |\mathcal{S}||\mathcal{D}|\big[T\delta^2(|\mathcal{S}|+1) + 2KT\delta^2|\mathcal{S}|\big]$ and $G_2 = (T-1)\delta(|\mathcal{S}|+1)/2$.

Proof. Squaring both sides of (4), rearranging, and bounding terms (see [22] for example) leads to

$$\frac{1}{2}\big[Q_i^d(t+T)^2 - Q_i^d(t)^2\big] \le Q_i^d(t)\left[\sum_{\tau=t}^{t+T-1} a_i^d(\tau) - Tb_i^d\right] + Q_i^d(t)\left[\sum_{j\in\mathcal{P}_i^d} x_{ji}^d(t,T) - \sum_{j\in\mathcal{H}_i^d} x_{ij}^d(t,T)\right] + C_i^d,$$

where $C_i^d = \frac{T^2\delta^2\left(|\mathcal{P}_i^d|+|\mathcal{H}_i^d|+2\right)}{2}$. Define a T-slot quadratic Lyapunov drift of queue backlogs [22] as

$$\Delta(t,T) = \frac{1}{2}\Big[\|Q(t+T)\|^2 - \|Q(t)\|^2\Big],$$

where $\|z\|$ is the l_2-norm of vector z, i.e., $\|Q(t)\|^2 = \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}} Q_i^d(t)^2$. It holds that

$$\Delta(t,T) \le \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}\left\{C_i^d + Q_i^d(t)\left[\sum_{\tau=t}^{t+T-1} a_i^d(\tau) - Tb_i^d\right]\right\} + \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}} Q_i^d(t)\left[\sum_{j\in\mathcal{P}_i^d} x_{ji}^d(t,T) - \sum_{j\in\mathcal{H}_i^d} x_{ij}^d(t,T)\right] \quad (12)$$

The second line of the above equation can be rewritten as

$$\sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}} Q_i^d(t)\left[\sum_{j\in\mathcal{P}_i^d} x_{ji}^d(t,T) - \sum_{j\in\mathcal{H}_i^d} x_{ij}^d(t,T)\right] = \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{H}_i}\sum_{d\in\mathcal{D}_{ij}} x_{ij}^d(t,T)\big[Q_j^d(t) - Q_i^d(t)\big], \quad (13)$$

using the fact that $x_{ij}^d(t,T) = 0$ for every $d \in \mathcal{D}\backslash\mathcal{D}_{ij}$.

Instead of minimizing the above expression, which leads to the MaxWeight algorithm, an state-dependent proximal term

$$\frac{k_{ij}(t)}{2}\left[x_{ij}^d(t,T) - \frac{x_{ij}^d(t-T,T)}{k_{ij}(t)}\right]^2 \quad \text{with}$$

$$k_{ij}(t) = \left[\frac{1}{Tc_{ij}}\sum_{d\in\mathcal{D}_{ij}}\big[Q_i^d(t) - Q_j^d(t) + x_{ij}^d(t-T,T)\big]_+\right]_{[1,K]}$$

is introduced[3], where $[z]_{[1,K]} = \max[1, \min[K, z]]$. This proximal term is non-negative and is upper bounded by $KT^2\delta^2$, so it holds from (12) and (13) that

$$\Delta(t,T) \le \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}\left\{C_i^d + Q_i^d(t)\left[\sum_{\tau=t}^{t+T-1} a_i^d(\tau) - Tb_i^d\right]\right\} + \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{H}_i}\sum_{d\in\mathcal{D}_{ij}}\left\{x_{ij}^d(t,T)\big[Q_j^d(t) - Q_i^d(t)\big] + \frac{k_{ij}(t)}{2}\left[x_{ij}^d(t,T) - \frac{x_{ij}^d(t-T,T)}{k_{ij}(t)}\right]^2\right\}. \quad (14)$$

[3]The $x_{ij}^d(t-T,T)$ in the proximal term can be replaced by $z_{ij}^d(t)/\alpha$ for any $0 \le z_{ij}^d(t) < \infty$ and $0 < \alpha < \infty$ to allow more control flexibility.

Minimizing the right-hand-side of (14) with respect to $\{x_{ij}^d(t,T)\}_{d\in\mathcal{D}_{ij}}$ leads to problem (5). Applying the result from Algorithm 1, which solves the minimization at reconfiguration time $t \in \mathcal{T}$, yields the bound for any other $\{\hat{x}_{ij}^d(t,T)\}_{d\in\mathcal{D}_{ij}}$ satisfying constraints in problem (5):

$$\Delta(t,T) \le \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}\left\{C_i^d + Q_i^d(t)\left[\sum_{\tau=t}^{t+T-1} a_i^d(\tau) - Tb_i^d\right]\right\} + \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{H}_i}\sum_{d\in\mathcal{D}_{ij}}\left\{\hat{x}_{ij}^d(t,T)\big[Q_j^d(t) - Q_i^d(t)\big] + \frac{k_{ij}(t)}{2}\left[\hat{x}_{ij}^d(t,T) - \frac{x_{ij}^d(t-T,T)}{k_{ij}(t)}\right]^2\right\}$$

Since the proximal term is bounded and policy $\left\{x_{ij}^{d*}(t,T) = \{x_{ij}^{d*}(\tau)\}_{\tau=t}^{t+T-1}\right\}_{d\in\mathcal{D}_{ij}}$, constructed from the randomized policy in Assumption 1, is one of those $\{\hat{x}_{ij}^d(t,T)\}_{d\in\mathcal{D}_{ij}}$, it follows that

$$\Delta(t,T) \le \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}\left\{C_i^d + Q_i^d(t)\left[\sum_{\tau=t}^{t+T-1} a_i^d(\tau) - Tb_i^d\right]\right\} + \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{H}_i}\sum_{d\in\mathcal{D}_{ij}}\left\{x_{ij}^{d*}(t,T)\big[Q_j^d(t) - Q_i^d(t)\big] + KT^2\delta^2\right\}$$

Applying identity (13), taking expectation, and using the independent property of the randomized policy gives

$$\mathbb{E}\left[\Delta(t,T)\right] \le \sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}\left\{D_i^d + \mathbb{E}\left[Q_i^d(t)\right]\times \sum_{\tau=t}^{t+T-1}\mathbb{E}\left[\sum_{j\in\mathcal{P}_i^d} x_{ji}^{d*}(\tau) + a_i^d(\tau) - \sum_{j\in\mathcal{H}_i^d} x_{ij}^{d*}(\tau) - b_i^d\right]\right\},$$

where $D_i^d = C_i^d + KT^2\delta^2\left(\left|\mathcal{P}_i^d\right| + \left|\mathcal{H}_i^d\right|\right)$. Assumption 1 implies for all $t \in \mathcal{T}$ that

$$\frac{1}{2T}\mathbb{E}\left[\|Q(t+T)\|^2 - \|Q(t)\|^2\right] \le G_1 - \epsilon\sum_{i\in\mathcal{S}}\sum_{d\in\mathcal{D}}\mathbb{E}\left[Q_i^d(t)\right], \quad (15)$$

where G_1 is defined in the theorem.

Queue dynamic (3) and the upper bound δ imply that $Q_i^d(t+\tau) \le Q_i^d(t) + \tau\delta(|\mathcal{S}|+1)$ for any $\tau \in \mathbb{Z}_+$ and $i \in \mathcal{S}, d \in \mathcal{D}$. Summing for $\tau \in \{0,1,\ldots,T-1\}$ gives $\sum_{\tau=0}^{T-1} Q_i^d(t+\tau) \le TQ_i^d(t) + T(T-1)\delta(|\mathcal{S}|+1)/2$ and

$$Q_i^d(t) \ge \frac{1}{T}\sum_{\tau=0}^{T-1} Q_i^d(t+\tau) - G_2 \quad \text{for } t \in \mathbb{Z}_+,$$

where G_2 is defined in the theorem. Substituting the above into inequality (15) yields:

$$\frac{1}{2T}\mathbb{E}\left[\|Q(t+T)\|^2 - \|Q(t)\|^2\right] \le G_1 + \epsilon G_2 - \frac{\epsilon}{T}\sum_{\tau=0}^{T-1}\mathbb{E}\left[\|Q(t+\tau)\|_1\right] \quad \text{for } t \in \mathcal{T}.$$

Telescope summation for $t \in \{0, T, \ldots, (N-1)T\}$ gives:

$$\frac{1}{2T}\mathbb{E}\left[\|Q(NT)\|^2 - \|Q(0)\|^2\right] \leq G_1 N + \epsilon G_2 N - \frac{\epsilon}{T}\sum_{n=0}^{N-1}\sum_{\tau=0}^{T-1}\mathbb{E}\left[\|Q(nT+\tau)\|_1\right].$$

Rearranging terms and taking supremum limit as $N \to \infty$ proves the theorem. □

IV. System Realization

The previous section provides an ideal allocation of decision variables. This section develops a heuristic improvement that is easier to implement for practical switches.

A. Approximation of Common Queues

An SDN switch has output queues at each of its ports [11], [29], [30]. Those queues can be assigned to each commodity. Let $Q^d_{ij}(t)$ denote the backlog of a queue for commodity d at the port of switch i connecting to switch j at time t. The queue backlog $Q^d_i(t)$ in Section II-D can be approximated by $\tilde{Q}^d_i(t) = \sum_{j\in\mathcal{H}^d_i} Q^d_{ij}(t)$, for $i \in \mathcal{S}, d \in \mathcal{D}$. It can be shown that this approximation becomes exact when the port's queues have never been emptied. Note that OpenFlow [30] allows 2^{32} unique queues per port, but availability of those queues may depend on a switch. This work encourages next generation switches to have a high number of available queues.

B. Additional Packet Headers

Three fields are appended into the IP header as IP options: `CommodityId`, `QueueInfo`, and `HashField`. The `CommodityId` identifies the commodity of `QueueInfo` which stores the rounded value of exponential moving average of approximated queue backlog, $\tilde{Q}^d_j(t)$. A packet from switch j to switch i carries queue information of a commodity, which is circularly selected from commodities in $\mathcal{D}_{ij}$[4]. The `HashField` is used for traffic splitting and is explained in Section IV-D.

Once a packet with the additional headers from switch j arrives to switch i, the contained queue information is extracted and stored in a local memory, which is denoted by $M^d_{ij}(t)$. This is the most recent queue information for commodity d on link (i,j) up to time t, where d is the value in `CommodityId`. The header processing can be implemented by P4 [32], NetFPGA [33], or a custom ASIC.

C. Weighted Fair Queueing

Each port of switch i connecting to switch $j \in \mathcal{H}_i$ is configured with weighted fair queueing. Let $r^d_{ij}(t,T)$ denote the measured number of commodity-d packets transmitted from switch i to switch j during the interval $[t, t+T)$. At every configuration time $t \in \mathcal{T}$, the weight $w^d_{ij}(t)$ for the interval $[t, t+T)$ is

$$w^d_{ij}(t) = \max\left[1, \tilde{Q}^d_i(t) - M^d_{ij}(t) + r^d_{ij}(t-T,T)/\alpha\right], d \in \mathcal{D}_{ij}$$

[4]This round-robin technique is inspired by CONGA [6]. The concept can also be applied to VXLAN [31].

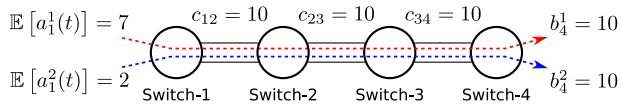

Fig. 7. Line Network with $\mathcal{S} = \{1,2,3,4\}, \mathcal{D} = \{1,2\}$ and $\mathcal{H}^d_1 = \{2\}, \mathcal{H}^d_2 = \{3\}, \mathcal{H}^d_3 = \{4\}$ for $d \in \mathcal{D}$

and $w^d_{ij}(t) = 0$ for $d \in \mathcal{D}\backslash\mathcal{D}_{ij}$. Parameter $\alpha > 0$ is added to scale the magnitude of actual traffic to match a finite queue capacity. From these weights, $w^d_{ij}(t)/\sum_{e\in\mathcal{D}} w^e_{ij}(t)$ fraction of link capacity is given to commodity d, which corresponds to the intuition in equation (6). Note that actual traffic $r^d_{ij}(t-T,T)$ is used instead of $x^d_{ij}(t-T,T)$, because it is a better approximation under TCP traffic.

D. Traffic Splitting by Hashing

A sending rate of a TCP connection is reduced when out-of-order packets are received at a destination. Hashing a packet to a next-hop switch based on `HashField` is implemented to reduce packet reordering. The `HashField` is generated once for the entire packet life. Packets from the same TCP connection have the same `HashField`, so they are hashed to the same path. Reordering does not occur if a hash rule at each switch is the same for the entire TCP connection. The hash rule is calculated as follows. For each commodity $d \in \mathcal{D}$, define $\left\{s^d_{ij}(t)\right\}_{j\in\mathcal{H}^d_i}$ as a solution of

$$\begin{aligned} \text{Maximize} \quad & \min_{j\in\mathcal{H}^d_i}\left\{Q^d_{ij}(t) - r^d_{ij}(t-T,T) + s^d_{ij}(t)\right\} \\ \text{Subject to} \quad & \sum_{j\in\mathcal{H}^d_i} s^d_{ij}(t) = \sum_{j\in\mathcal{H}^d_i} r^d_{ij}(t-T,T) \\ & s^d_{ij}(t) \in \mathbb{Z}_+ \quad \text{for all } j \in \mathcal{H}^d_i. \end{aligned} \tag{16}$$

This problem can be solved in polynomial time. Let $s^d_{ij}(t) = 0$ for all $j \in \mathcal{H}^d_i$. Iteratively, $s^d_{ij}(t)$ is increased for a group of indices in $\text{Argmin}_{j\in\mathcal{H}^d_i}\left[Q^d_{ij}(t) - r^d_{ij}(t-T,T) + s^d_{ij}(t)\right]$ until the equality constraint is met[5]. This process keeps increasing the $\min_{j\in\mathcal{H}^d_i}\left\{Q^d_{ij}(t) - r^d_{ij}(t-T,T) + s^d_{ij}(t)\right\}$. The intuition of problem (16) is that it attempts to equalize the backlog levels at all ports at end of the interval $[t, t+T)$ by using $r^d_{ij}(t-T,T)$ as an estimate of actual transmission $r^d_{ij}(t,T)$. The attempt tries to make the approximation in Section IV-A exact. The splitting ratio of the port connecting to switch j is $f^d_{ij}(t) = s^d_{ij}(t)/\sum_{k\in\mathcal{H}^d_i} s^d_{ik}(t)$ for $j \in \mathcal{H}^d_i$.

V. Simulations

A. Ideal Simulation

Algorithm 1 is simulated according to the system model in Section II. A switch uses a token bucket mechanism to ensure that transmission rate per interval satisfy constraint (1) after $x^d_{ij}(t,T)$ is determined.

A line network in Fig. 7 is simulated with the interval length $T = 100$ and the constant $K = 10$. The network is simulated for 10^5 slots. The average backlogs shown in Table

[5]The process can be accelerated by increasing multiple values of $s^d_{ij}(t)$'s per iteration.

TABLE I
AVERAGE BACKLOGS UNDER ALGORITHM 1 AND MAXWEIGHT

	Commodity 1		Commodity 2	
	Algorithm 1	MaxWeight	Algorithm 1	MaxWeight
Switch-1	14.88	2598.19	6.25	213.45
Switch-2	8.17	1699.99	2.35	195.12
Switch-3	7.52	700.01	2.16	199.51
Switch-4	7.12	7.00	2.00	2.00

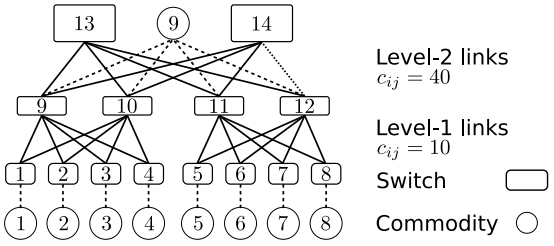

Fig. 8. Intra datacenter network with $\mathcal{S} = \{1, 2, \ldots, 14\}, \mathcal{D} = \{1, 2, \ldots, 9\}$. Each next-hop set $\mathcal{H}_i^d$ contains next-hop switches with the shortest distance to commodity d, e.g., $\mathcal{H}_1^8 = \{9, 10\}, \mathcal{H}_9^8 = \{13, 14\}, \mathcal{H}_{13}^8 = \{11, 12\}, \mathcal{H}_{11}^8 = \{8\}, \mathcal{H}_1^9 = \{9, 10\} = \mathcal{H}_1^e$ for $e \in \{2, 3, 4\}$. Arrivals are $\mathbb{E}\left[a_i^d(t)\right] = 2$ and $\mathbb{E}\left[a_i^9(t)\right] = 1$ for $d, i \in \{1, \ldots, 8\}$; otherwise 0. Departure rate is $b_i^d = 20$ if commodity d connects to switch i; otherwise 0.

I are calculated after Algorithm 1 and MaxWeight converge. The link-capacity sharing of Algorithm 1 leads to small queue backlogs. For scaling comparison, when $T = 1000$, the average backlogs of commodity 1 at switch 1 are respectively 23.94 and 16005.63 under Algorithm 1 and MaxWeight. In this simulation, the maximum value of $k_{ij}(t)$ for all $i \in \mathcal{S}, j \in \mathcal{H}_i, t \in \{0, \ldots, 10^5\}$ is $1.732 < K$.

A network in Fig. 8 is simulated with $T = 100$ and $K = 10$. After 10^5 slots, the average backlogs per queue (from all 124 queues) are 34.23 and 555.95 under Algorithm 1 and MaxWeight. An event $k_{ij}(t) = K$ occurs 93.41% of the times due to that the network operates near its capacity boundary. Reducing the arrivals by 12% (24%) yields 20.68% (5.49%) of the times that $k_{ij}(t) = K$. This suggests that Algorithm 2 is rarely invoked, i.e., $k_{ij}(t) < K$, when a network does not operate near its capacity boundary. In practice, TCP flows with congestion control are different from the i.i.d. arrivals, so Algorithm 2 is not included in the heuristic algorithm.

B. Network Simulator

The heuristic in-network load-balancing algorithm in Section IV is simulated by OMNeT++ [28]. All simulations share the following setting. Capacity of each commodity queue at a switch port is 200 packets. For ECMP setting, a shared queue at a switch port has buffer capacity of $200 \times |\mathcal{D}|$ packets, where $\mathcal{D}$ is a set of commodities in a considered network. Configuration interval is $T = 1$ms, and the scaling parameter is $\alpha = 5$. The NewReno TCP from INET Framework [34] is adjusted for 10Gbps and 40Gbps link speeds. Every TCP flow is randomly established during $[0s, 0.5s]$ and starts during $[1s, 1.01s]$. Each flow transmits 1MB of data. Flow completion time (FCT) is measured as the performance metrics, which are also used in [6], [11]. FCT is the duration of time to complete a flow, i.e., the time to send 1MB of data.

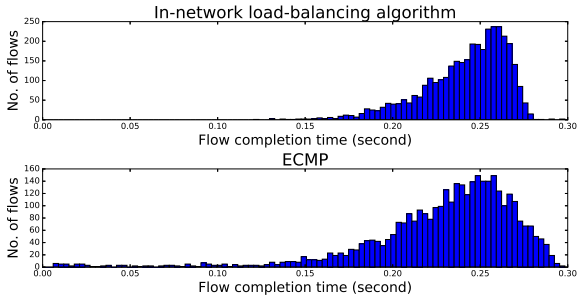

Fig. 9. The FCTs from the network in Fig. 8 without commodity 9

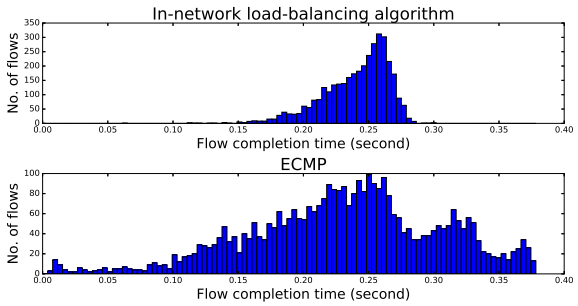

Fig. 10. The FCTs of all flows in the network in Fig. 8 where commodity 9 is omitted and the link between switches 12 and 14 fails

A network without commodity 9 in Fig. 8 is simulated. The speeds of level-1 links and level-2 links are respectively 10Gbps and 40Gbps. 64 flows are generated from a commodity to one another, and the total number of flows in the network is 3584. The FCTs under the heuristic algorithm and ECMP are shown in Fig. 9, and their variances are 5.8×10^{-4} and 19.5×10^{-4}. Unsurprisingly, they are comparable, since the topology is optimized for ECMP. The FCTs under the heuristic algorithm has less variation, as the distribution of flows are more balance. Note that the tail of FCTs is critical for interactive services [11].

A follow-up scenario is simulated when the link between switches 12 and 14 in Fig. 8 fails. The FCTs of all flows are shown Fig. 10. The heuristic algorithm balances the flows better than ECMP. Switches 9 and 10 hash more flows to switch 13 than switch 14, while ECMP hashes flows equally.

A network in Fig. 11 illustrates the adaptiveness of the heuristic algorithm when some link capacity is taken away by priority flows. The FCTs of all flows are shown in Fig. 12. The FCTs under the heuristic algorithm is more balance compared to the FCTs under ECMP, as switches 1 and 2 hash more flows to switch 4 instead of equally hashing in ECMP case. Note that if the priority flows begin shortly after 1.01s, the same result is observed.

A similar trend can be observed from scenarios with short

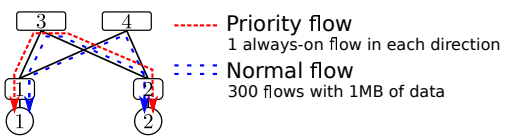

Fig. 11. A network with a priority flow in each direction and 600 normal flows between commodities 1 and 2

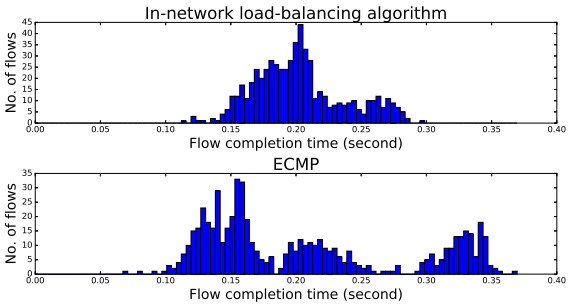

Fig. 12. The FCTs of normal flows between commodities 1 and 2 in Fig. 11

flows (10KB of data per flow) when $T = 0.1$ms and $\alpha = 1$.

VI. Conclusion

This paper showed that practical load-balancing in datacenter networks can be tackled by the concept of throughput optimality. Our first algorithm is a novel variation of the MaxWeight concept that treats control plane and data plane timescales (useful for software defined networking), allows link capacity sharing during a control plane interval, incorporates weighted fair queueing aspects, and comes with a proof of throughput optimality. Next, this algorithm was modified to include heuristic improvements that allow easy operation with practical switch capabilities and works gracefully with TCP flows. Ideal and OMNeT++ simulations show promising potential against existing MaxWeight and ECMP.

References

[1] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta, "Vl2: A scalable and flexible data center network," *SIGCOMM Comput. Commun. Rev.*, Aug. 2009.

[2] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," *SIGCOMM Comput. Commun. Rev.*, Aug. 2008.

[3] A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano, A. Kanagala, J. Provost, J. Simmons, E. Tanda, J. Wanderer, U. Hölzle, S. Stuart, and A. Vahdat, "Jupiter rising: A decade of clos topologies and centralized control in google's datacenter network," *SIGCOMM Comput. Commun. Rev.*, Aug. 2015.

[4] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren, "Inside the social network's (datacenter) network," *SIGCOMM Comput. Commun. Rev.*, Aug. 2015.

[5] N. Kang, M. Ghobadi, J. Reumann, A. Shraer, and J. Rexford, "Efficient traffic splitting on commodity switches," in *Proceedings of the 11th ACM International on Conference on emerging Networking Experiments and Technologies (CoNEXT)*, 2015.

[6] M. Alizadeh, T. Edsall, S. Dharmapurikar, R. Vaidyanathan, K. Chu, A. Fingerhut, V. T. Lam, F. Matus, R. Pan, N. Yadav, and G. Varghese, "Conga: Distributed congestion-aware load balancing for datacenters," *SIGCOMM Comput. Commun. Rev.*, Aug. 2014.

[7] J. Zhou, M. Tewari, M. Zhu, A. Kabbani, L. Poutievski, A. Singh, and A. Vahdat, "Wcmp: Weighted cost multipathing for improved fairness in data centers," in *Proceedings of the Ninth European Conference on Computer Systems*, ser. EuroSys '14, 2014.

[8] D. L. H. W. X. Z. Y. S. C. T. Y. Z. S. L. G. L. Chuanxiong Guo, Guohan Lu, "Bcube: A high performance, server-centric network architecture for modular data centers," Aug. 2009.

[9] Software-defined networking: The new norm for networks. [Online]. Available: https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf

[10] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: Dynamic flow scheduling for data center networks," in *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'10, 2010.

[11] D. Zats, T. Das, P. Mohan, D. Borthakur, and R. Katz, "Detail: Reducing the flow completion time tail in datacenter networks," *SIGCOMM Comput. Commun. Rev.*, Aug. 2012.

[12] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *Automatic Control, IEEE Transactions on*, Dec. 1992.

[13] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Transactions on Communications*, Aug. 1999.

[14] T. Javidi, C. H. Wang, and T. Akta, "A novel data center network architecture with zero in-network queuing," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt), 2015 13th International Symposium on*, May. 2015.

[15] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE Journal on Selected Areas in Communications*, Jan. 2005.

[16] A. Eryilmaz and R. Srikant, "Joint congestion control, routing, and mac for stability and fairness in wireless networks," *IEEE Journal on Selected Areas in Communications*, Aug. 2006.

[17] S. Supittayapornpong and M. Neely, "Achieving utility-delay-reliability tradeoff in stochastic network optimization with finite buffers," in *Computer Communications (INFOCOM), 2015 IEEE Conference on*, Apr. 2015.

[18] L. B. Le, E. Modiano, and N. B. Shroff, "Optimal control of wireless networks with finite buffers," *IEEE/ACM Transactions on Networking*, Aug. 2012.

[19] H. Seferoglu and E. Modiano, "Tcp-aware backpressure routing and scheduling," *IEEE Transactions on Mobile Computing*, Jul. 2016.

[20] S. Moeller, A. Sridharan, B. Krishnamachari, and O. Gnawali, "Routing without routes: The backpressure collection protocol," in *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, ser. IPSN '10, 2010.

[21] L. Bui, R. Srikant, and A. Stolyar, "Novel architectures and algorithms for delay reduction in back-pressure scheduling and routing," in *Computer Communications (INFOCOM), 2009 IEEE Conference on*, Apr. 2009.

[22] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, 2010.

[23] S. Supittayapornpong and M. Neely, "Quality of information maximization for wireless networks via a fully separable quadratic policy," *Networking, IEEE/ACM Transactions on*, Apr. 2015.

[24] L. Huang and M. J. Neely, "Delay reduction via lagrange multipliers in stochastic network optimization," *Automatic Control, IEEE Transactions on*, Apr. 2011.

[25] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," *SIGCOMM Comput. Commun. Rev.*, Aug. 1989.

[26] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the single-node case," *IEEE/ACM Transactions on Networking*, Jun. 1993.

[27] M. Shafiee and J. Ghaderi, "A simple congestion-aware algorithm for load balancing in datacenter networks," in *Computer Communications (INFOCOM), 2016 IEEE Conference on*, Apr. 2016.

[28] Omnet++. [Online]. Available: https://omnetpp.org/

[29] H. Seferoglu and E. Modiano, "Separation of routing and scheduling in backpressure-based wireless networks," *IEEE/ACM Transactions on Networking*, Jun. 2016.

[30] Openflow switch specification: Version 1.5.1. [Online]. Available: https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-switch-v1.5.1.pdf

[31] Virtual extensible local area network (vxlan). [Online]. Available: https://tools.ietf.org/html/rfc7348

[32] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: Programming protocol-independent packet processors," *SIGCOMM Comput. Commun. Rev.*, Jul. 2014.

[33] J. W. Lockwood, N. McKeown, G. Watson, G. Gibb, P. Hartke, J. Naous, R. Raghuraman, and J. Luo, "Netfpga–an open platform for gigabit-rate network switching and routing," in *2007 IEEE International Conference on Microelectronic Systems Education (MSE'07)*, Jun. 2007.

[34] INET framework. [Online]. Available: https://inet.omnetpp.org/